

# Package ‘clustermq’

April 22, 2018

**Title** Evaluate Function Calls on HPC Schedulers (LSF, SGE, SLURM, PBS/Torque)

**Version** 0.8.4

**Author** Michael Schubert <mschu.dev@gmail.com>

**Maintainer** Michael Schubert <mschu.dev@gmail.com>

**Description** Evaluate arbitrary function calls using workers on HPC schedulers in single line of code. All processing is done on the network without accessing the file system. Remote schedulers are supported via SSH.

**URL** <https://github.com/mschubert/clustermq>

**BugReports** <https://github.com/mschubert/clustermq/issues>

**Depends** R (>= 3.0.2)

**Imports** infuser (>= 0.2.8), narray, progress, purrr, R6, rzmq (>= 0.9.3), utils

**License** Apache License (== 2.0) | file LICENSE

**LazyData** true

**Encoding** UTF-8

**Suggests** dplyr, knitr, parallel, roxygen2 (>= 5.0.0), testthat, tools

**VignetteBuilder** knitr

**RoxygenNote** 6.0.1.9000

**NeedsCompilation** no

**Repository** CRAN

**Date/Publication** 2018-04-22 17:24:28 UTC

## R topics documented:

bind_avail . . . . .	2
check_args . . . . .	3
chunk . . . . .	3
clustermq . . . . .	4

DO_CHUNK . . . . .	4
DO_SETUP . . . . .	5
LOCAL . . . . .	5
LSF . . . . .	5
master . . . . .	6
MULTICORE . . . . .	6
PBS . . . . .	7
PROXY_CMD . . . . .	7
PROXY_READY . . . . .	7
PROXY_STOP . . . . .	8
PROXY_UP . . . . .	8
purrr_lookup . . . . .	8
Q . . . . .	9
QSys . . . . .	10
Q_rows . . . . .	10
SGE . . . . .	11
SLURM . . . . .	12
SSH . . . . .	12
ssh_proxy . . . . .	12
summarize_result . . . . .	13
vec_lookup . . . . .	13
worker . . . . .	14
workers . . . . .	14
WORKER_DONE . . . . .	15
WORKER_READY . . . . .	15
WORKER_STOP . . . . .	15
WORKER_UP . . . . .	16
work_chunk . . . . .	16
<b>Index</b>	<b>17</b>

---

bind_avail	<i>Binds an rzmq to an available port in given range</i>
------------	--

---

## Description

Binds an rzmq to an available port in given range

## Usage

```
bind_avail(socket, range, iface = "tcp://*", n_tries = 100)
```

## Arguments

socket	An rzmq socket object
range	Numbers to consider (e.g. 6000:8000)
iface	Interface to listen on
n_tries	Number of ports to try in range

**Value**

The port the socket is bound to

---

check_args	<i>Function to check arguments with which Q() is called</i>
------------	---

---

**Description**

Function to check arguments with which Q() is called

**Usage**

```
check_args(fun, iter, const = list())
```

**Arguments**

fun	A function to call
iter	Objects to be iterated in each function call
const	A list of constant arguments passed to each function call

**Value**

Processed iterated argument list if 'iter' is a list

---

chunk	<i>Subset index chunk for processing</i>
-------	--

---

**Description**

'attr' in '[.data.frame' takes too much CPU time

**Usage**

```
chunk(x, i)
```

**Arguments**

x	Index data.frame
i	Rows to subset

**Value**

x[i,]

clustermq

*Evaluate Function Calls on HPC Schedulers (LSF, SGE, SLURM)*

---

**Description**

Provides the Q function to send arbitrary function calls to workers on HPC schedulers without relying on network-mounted storage. Allows using remote schedulers via SSH.

**Details**

Under the hood, this will submit a cluster job that connects to the master via TCP the master will then send the function and argument chunks to the worker and the worker will return the results to the master until everything is done and you get back your result

Computations are done entirely on the network and without any temporary files on network-mounted storage, so there is no strain on the file system apart from starting up R once per job. This removes the biggest bottleneck in distributed computing.

Using this approach, we can easily do load-balancing, i.e. workers that get their jobs done faster will also receive more function calls to work on. This is especially useful if not all calls return after the same time, or one worker has a high load.

For more detailed usage instructions, see the documentation of the Q function.

---

DO\_CHUNK*Chunk of iterated arguments for the worker*

---

**Description**

Field has to be 'chunk'

**Usage**

DO\_CHUNK

**Format**

An object of class integer of length 1.

---

DO_SETUP	<i>Message contains common data</i>
----------	-------------------------------------

---

**Description**

This includes the function definition, common data, export

**Usage**

DO\_SETUP

**Format**

An object of class integer of length 1.

---

LOCAL	<i>Placeholder for local processing</i>
-------	---

---

**Description**

Mainly so tests pass without setting up a scheduler

**Usage**

LOCAL

**Format**

An object of class R6ClassGenerator of length 24.

---

LSF	<i>LSF scheduler functions</i>
-----	--------------------------------

---

**Description**

Derives from QSys to provide LSF-specific functions

**Usage**

LSF

**Format**

An object of class R6ClassGenerator of length 26.

---

master                      *Master controlling the workers*

---

### Description

exchanging messages between the master and workers works the following way: \* we have submitted a job where we don't know when it will start up \* it starts, sends is a message list(id=0) indicating it is ready \* we send it the function definition and common data \* we also send it the first data set to work on \* when we get any id > 0, it is a result that we store \* and send the next data set/index to work on \* when computatons are complete, we send id=0 to the worker \* it responds with id=-1 (and usage stats) and shuts down

### Usage

```
master(qsys, iter, rettype = "list", fail_on_error = TRUE, wait_time = NA,
      chunk_size = NA, timeout = Inf)
```

### Arguments

qsys	Instance of QSys object
iter	Objects to be iterated in each function call
rettype	Return type of function
fail_on_error	If an error occurs on the workers, continue or fail?
wait_time	Time to wait between messages; set 0 for short calls defaults to 1/sqrt(number_of_functon_calls)
chunk_size	Number of function calls to chunk together defaults to 100 chunks per worker or max. 500 kb per chunk
timeout	Maximum time in seconds to wait for worker (default: Inf)

### Value

A list of whatever 'fun' returned

---

MULTICORE                      *Process on multiple cores on one machine*

---

### Description

This makes use of rzmq messaging and sends requests via TCP/IP

### Usage

```
MULTICORE
```

### Format

An object of class R6ClassGenerator of length 24.

---

PBS	<i>PBS/Torque scheduler functions</i>
-----	---------------------------------------

---

**Description**

Derives from QSys to provide PBS/Torque-specific functions

**Usage**

PBS

**Format**

An object of class R6ClassGenerator of length 26.

---

PROXY_CMD	<i>Message is an SSH command</i>
-----------	----------------------------------

---

**Description**

Field is either 'exec' (command to run) or 'reply' (how it went)

**Usage**

PROXY\_CMD

**Format**

An object of class integer of length 1.

---

PROXY_READY	<i>Message ID indicating SSH proxy is ready to distribute data</i>
-------------	--

---

**Description**

Field has to be 'proxy'

**Usage**

PROXY\_READY

**Format**

An object of class integer of length 1.

---

PROXY_STOP	<i>Message telling the SSH proxy to clean up</i>
------------	--

---

**Description**

No fields. Signals the worker to break its main loop.

**Usage**

PROXY\_STOP

**Format**

An object of class integer of length 1.

---

PROXY_UP	<i>Message ID indicating SSH proxy is up</i>
----------	--

---

**Description**

Message ID indicating SSH proxy is up

**Usage**

PROXY\_UP

**Format**

An object of class integer of length 1.

---

purrr_lookup	<i>Lookup table for return types to purrr functions</i>
--------------	---

---

**Description**

Lookup table for return types to purrr functions

**Usage**

purrr\_lookup

**Format**

An object of class list of length 9.



**Description**

Queue function calls on the cluster

**Usage**

```
Q(fun, ..., const = list(), export = list(), seed = 128965,
  memory = NULL, template = list(), n_jobs = NULL, job_size = NULL,
  split_array_by = -1, rettype = "list", fail_on_error = TRUE,
  workers = NULL, log_worker = FALSE, wait_time = NA, chunk_size = NA,
  timeout = Inf)
```

**Arguments**

fun	A function to call
...	Objects to be iterated in each function call
const	A list of constant arguments passed to each function call
export	List of objects to be exported to the worker
seed	A seed to set for each function call
memory	Short for template=list(memory=value)
template	A named list of values to fill in template
n_jobs	The number of LSF jobs to submit; upper limit of jobs if job_size is given as well
job_size	The number of function calls per job
split_array_by	The dimension number to split any arrays in '...'; default: last
rettype	Return type of function call (vector type or 'list')
fail_on_error	If an error occurs on the workers, continue or fail?
workers	Optional instance of QSys representing a worker pool
log_worker	Write a log file for each worker
wait_time	Time to wait between messages; set 0 for short calls defaults to 1/sqrt(number_of_function_calls)
chunk_size	Number of function calls to chunk together defaults to 100 chunks per worker or max. 10 kb per chunk
timeout	Maximum time in seconds to wait for worker (default: Inf)

**Value**

A list of whatever 'fun' returned

**Examples**

```
## Not run:
# Run a simple multiplication for numbers 1 to 3 on a worker node
fx = function(x) x * 2
Q(fx, x=1:3, n_jobs=1)
# list(2,4,6)

# Run a mutate() call in dplyr on a worker node
iris %>%
  mutate(area = Q(`*`, e1=Sepal.Length, e2=Sepal.Width, n_jobs=1))
# iris with an additional column 'area'

## End(Not run)
```

---

QSys

*Class for basic queuing system functions*


---

**Description**

Provides the basic functions needed to communicate between machines This should abstract most functions of rZMQ so the scheduler implementations can rely on the higher level functionality

**Usage**

QSys

**Format**

An object of class R6ClassGenerator of length 24.

---

Q\_rows

*Queue function calls defined by rows in a data.frame*


---

**Description**

Queue function calls defined by rows in a data.frame

**Usage**

```
Q_rows(df, fun, const = list(), export = list(), seed = 128965,
memory = NULL, template = list(), n_jobs = NULL, job_size = NULL,
rettype = "list", fail_on_error = TRUE, workers = NULL,
log_worker = FALSE, wait_time = NA, chunk_size = NA, timeout = Inf)
```

**Arguments**

df	data.frame with iterated arguments
fun	A function to call
const	A list of constant arguments passed to each function call
export	List of objects to be exported to the worker
seed	A seed to set for each function call
memory	Short for template=list(memory=value)
template	A named list of values to fill in template
n_jobs	The number of LSF jobs to submit; upper limit of jobs if job_size is given as well
job_size	The number of function calls per job
rettype	Return type of function call (vector type or 'list')
fail_on_error	If an error occurs on the workers, continue or fail?
workers	Optional instance of QSys representing a worker pool
log_worker	Write a log file for each worker
wait_time	Time to wait between messages; set 0 for short calls defaults to 1/sqrt(number_of_function_calls)
chunk_size	Number of function calls to chunk together defaults to 100 chunks per worker or max. 10 kb per chunk
timeout	Maximum time in seconds to wait for worker (default: Inf)

---

 SGE

---

*SGE scheduler functions*


---

**Description**

Derives from QSys to provide SGE-specific functions

**Usage**

SGE

**Format**

An object of class R6ClassGenerator of length 26.

---

SLURM	<i>SLURM scheduler functions</i>
-------	----------------------------------

---

**Description**

Derives from QSys to provide SLURM-specific functions

**Usage**

SLURM

**Format**

An object of class R6ClassGenerator of length 26.

---

SSH	<i>SSH scheduler functions</i>
-----	--------------------------------

---

**Description**

Derives from QSys to provide SSH-specific functions

**Usage**

SSH

**Format**

An object of class R6ClassGenerator of length 25.

---

ssh_proxy	<i>SSH proxy for different schedulers</i>
-----------	---

---

**Description**

Do not call this manually, the SSH qsys will do that

**Usage**

```
ssh_proxy(ctl, job, qsys_id = qsys_default)
```

**Arguments**

ctl	The port to connect to the master for proxy control
job	The port to connect to the master for job control
qsys_id	Character string of QSys class to use

---

summarize_result	<i>Print a summary of errors and warnings that occurred during processing</i>
------------------	---

---

**Description**

Print a summary of errors and warnings that occurred during processing

**Usage**

```
summarize_result(result, n_errors, n_warnings, cond_msgs, at = length(result),
  fail_on_error = TRUE)
```

**Arguments**

result	A list or vector of the processing result
n_errors	How many errors occurred
n_warnings	How many warnings occurred
cond_msgs	Error and warnings messages, we display first 50
at	How many calls were procesed up to this point
fail_on_error	Stop if error(s) occurred

---

vec_lookup	<i>Lookup table for return types to vector NAs</i>
------------	--

---

**Description**

Lookup table for return types to vector NAs

**Usage**

```
vec_lookup
```

**Format**

An object of class `list` of length 9.

---

worker	<i>R worker submitted as cluster job</i>
--------	--

---

**Description**

Do not call this manually, the master will do that

**Usage**

```
worker(master, timeout = 600, ..., verbose = TRUE)
```

**Arguments**

master	The master address (tcp://ip:port)
timeout	Time until worker shuts down without hearing from master
...	Catch-all to not break older template values (ignored)
verbose	Whether to print debug messages

---

workers	<i>Creates a pool of workers</i>
---------	----------------------------------

---

**Description**

Creates a pool of workers

**Usage**

```
workers(n_jobs, data = NULL, reuse = TRUE, template = list(),
        log_worker = FALSE, qsys_id = qsys_default)
```

**Arguments**

n_jobs	Number of jobs to submit (0 implies local processing)
data	Set common data (function, constant args, seed)
reuse	Whether workers are reusable or get shut down after call
template	A named list of values to fill in template
log_worker	Write a log file for each worker
qsys_id	Character string of QSys class to use

**Value**

An instance of the QSys class

---

WORKER_DONE	<i>Message ID indicating worker is shutting down</i>
-------------	--

---

**Description**

Field has to be 'time' with a object returned by 'Sys.time()'

**Usage**

WORKER\_DONE

**Format**

An object of class integer of length 1.

---

WORKER_READY	<i>Message ID indicating worker is accepting jobs</i>
--------------	---

---

**Description**

It may contain the field 'result' with a finished chunk

**Usage**

WORKER\_READY

**Format**

An object of class integer of length 1.

---

WORKER_STOP	<i>Message ID telling worker to stop</i>
-------------	--

---

**Description**

No fields

**Usage**

WORKER\_STOP

**Format**

An object of class integer of length 1.

---

WORKER_UP	<i>Message ID indicating worker is accepting jobs</i>
-----------	---

---

**Description**

Field has to be 'worker\_id' to master or empty to ssh\_proxy Answer is serialized common data ('fun', 'const', and 'seed') or 'redirect' (with URL where worker can get data)

**Usage**

WORKER\_UP

**Format**

An object of class integer of length 1.

---

work_chunk	<i>Function to process a chunk of calls</i>
------------	---

---

**Description**

Each chunk comes encapsulated in a data.frame

**Usage**

```
work_chunk(df, fun, const_args = list(), rettype = "list",
           common_seed = NULL, progress = FALSE)
```

**Arguments**

df	A data.frame with call IDs as rownames and arguments as columns
fun	The function to call
const_args	Constant arguments passed to each call
rettype	Return type of function
common_seed	A seed offset common to all function calls
progress	Logical indicated whether to display a progress bar

**Value**

A list of call results (or try-error if they failed)



# Index

## \*Topic **datasets**

- DO\_CHUNK, 4
  - DO\_SETUP, 5
  - LOCAL, 5
  - LSF, 5
  - MULTICORE, 6
  - PBS, 7
  - PROXY\_CMD, 7
  - PROXY\_READY, 7
  - PROXY\_STOP, 8
  - PROXY\_UP, 8
  - purrr\_lookup, 8
  - QSys, 10
  - SGE, 11
  - SLURM, 12
  - SSH, 12
  - vec\_lookup, 13
  - WORKER\_DONE, 15
  - WORKER\_READY, 15
  - WORKER\_STOP, 15
  - WORKER\_UP, 16
- bind\_avail, 2
- check\_args, 3
- chunk, 3
- clustermq, 4
- clustermq-package (clustermq), 4
- DO\_CHUNK, 4
- DO\_SETUP, 5
- LOCAL, 5
- LSF, 5
- master, 6
- MULTICORE, 6
- PBS, 7
- PROXY\_CMD, 7
- PROXY\_READY, 7
- PROXY\_STOP, 8
- PROXY\_UP, 8
- purrr\_lookup, 8
- Q, 9
- Q\_rows, 10
- QSys, 10
- SGE, 11
- SLURM, 12
- SSH, 12
- ssh\_proxy, 12
- summarize\_result, 13
- vec\_lookup, 13
- work\_chunk, 16
- worker, 14
- WORKER\_DONE, 15
- WORKER\_READY, 15
- WORKER\_STOP, 15
- WORKER\_UP, 16
- workers, 14